

CDF Data Guard Briefing

Nelly Stanfield

CSS/DSG

November 13, 2003

CDF Goals

- return to data taking within 1 hour after a hard failure and CDF experts on site
- a form of the database remains operational for reading and permanent writing (data ultimately saved as part of production database) during long downtimes (several hours of planned switchover time acceptable).

CDF Requirements

- Mixed Replication(Basic/Streams) and Data Guard environment
- No data loss during downtime – required
- All aspects supported by Oracle
- Do not increase sysadmin or DBA load unreasonably
- No more than 1 run lost (4 hours data=downtime) at hard failure

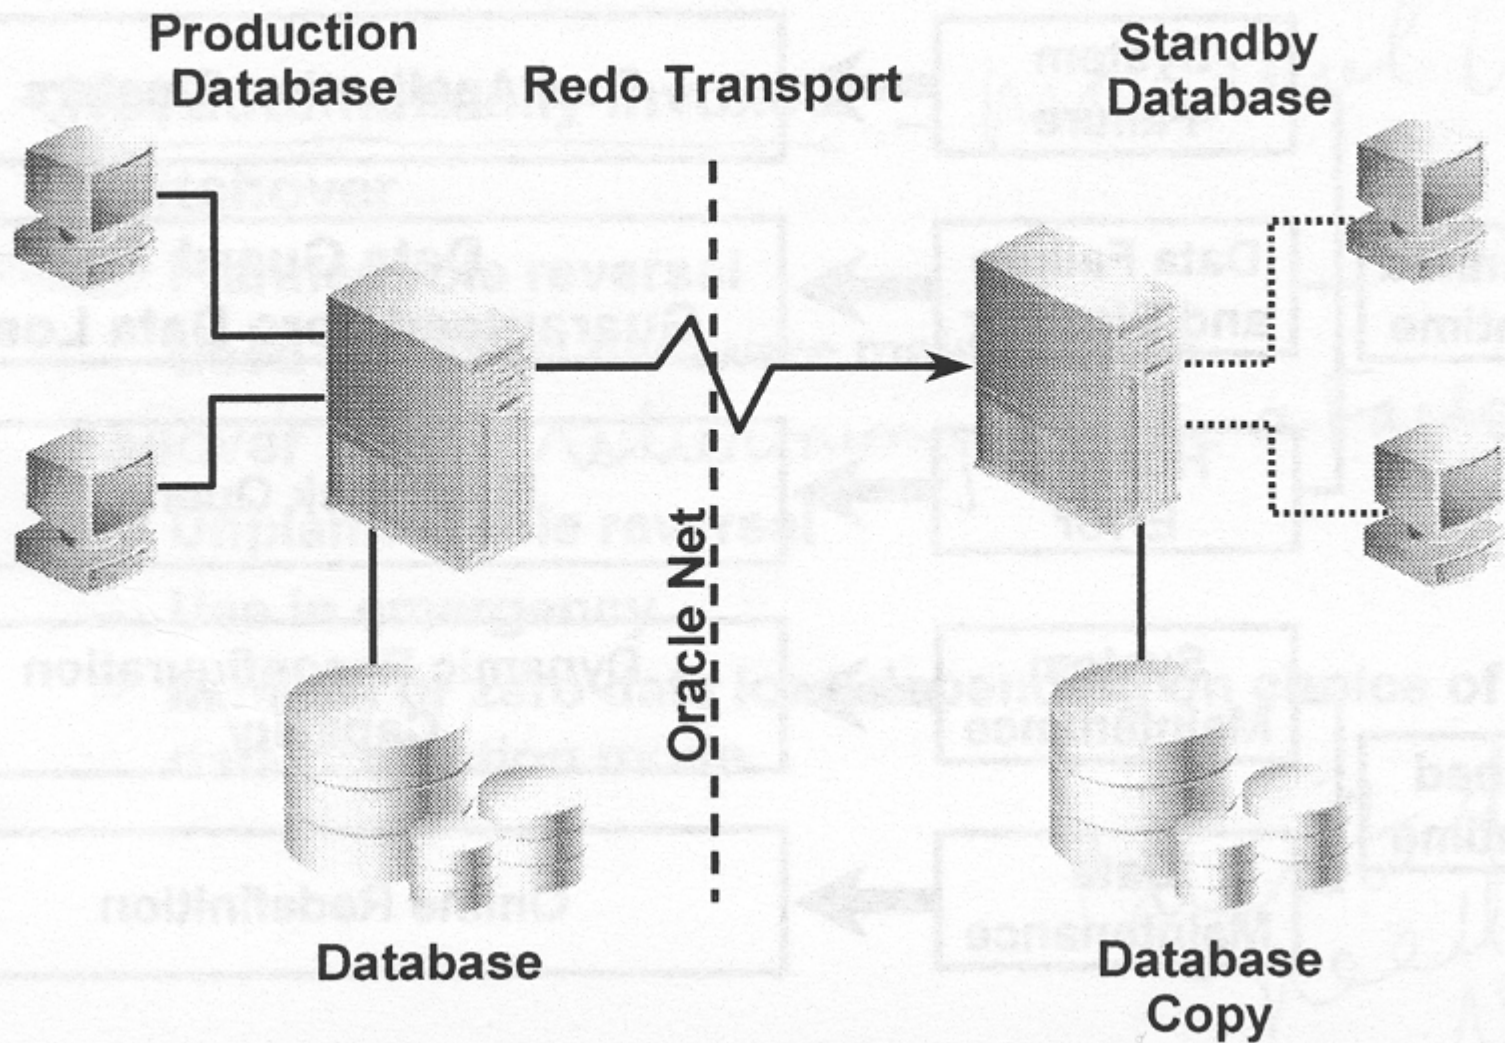
CDF Requirements (con't)

- Uses linux platforms that will not need replacing before the end of Run II.
- In place in less than 6 months
- < \$70K hardware
- Zero data loss at hard failure
- Transparent (no intervention) failover

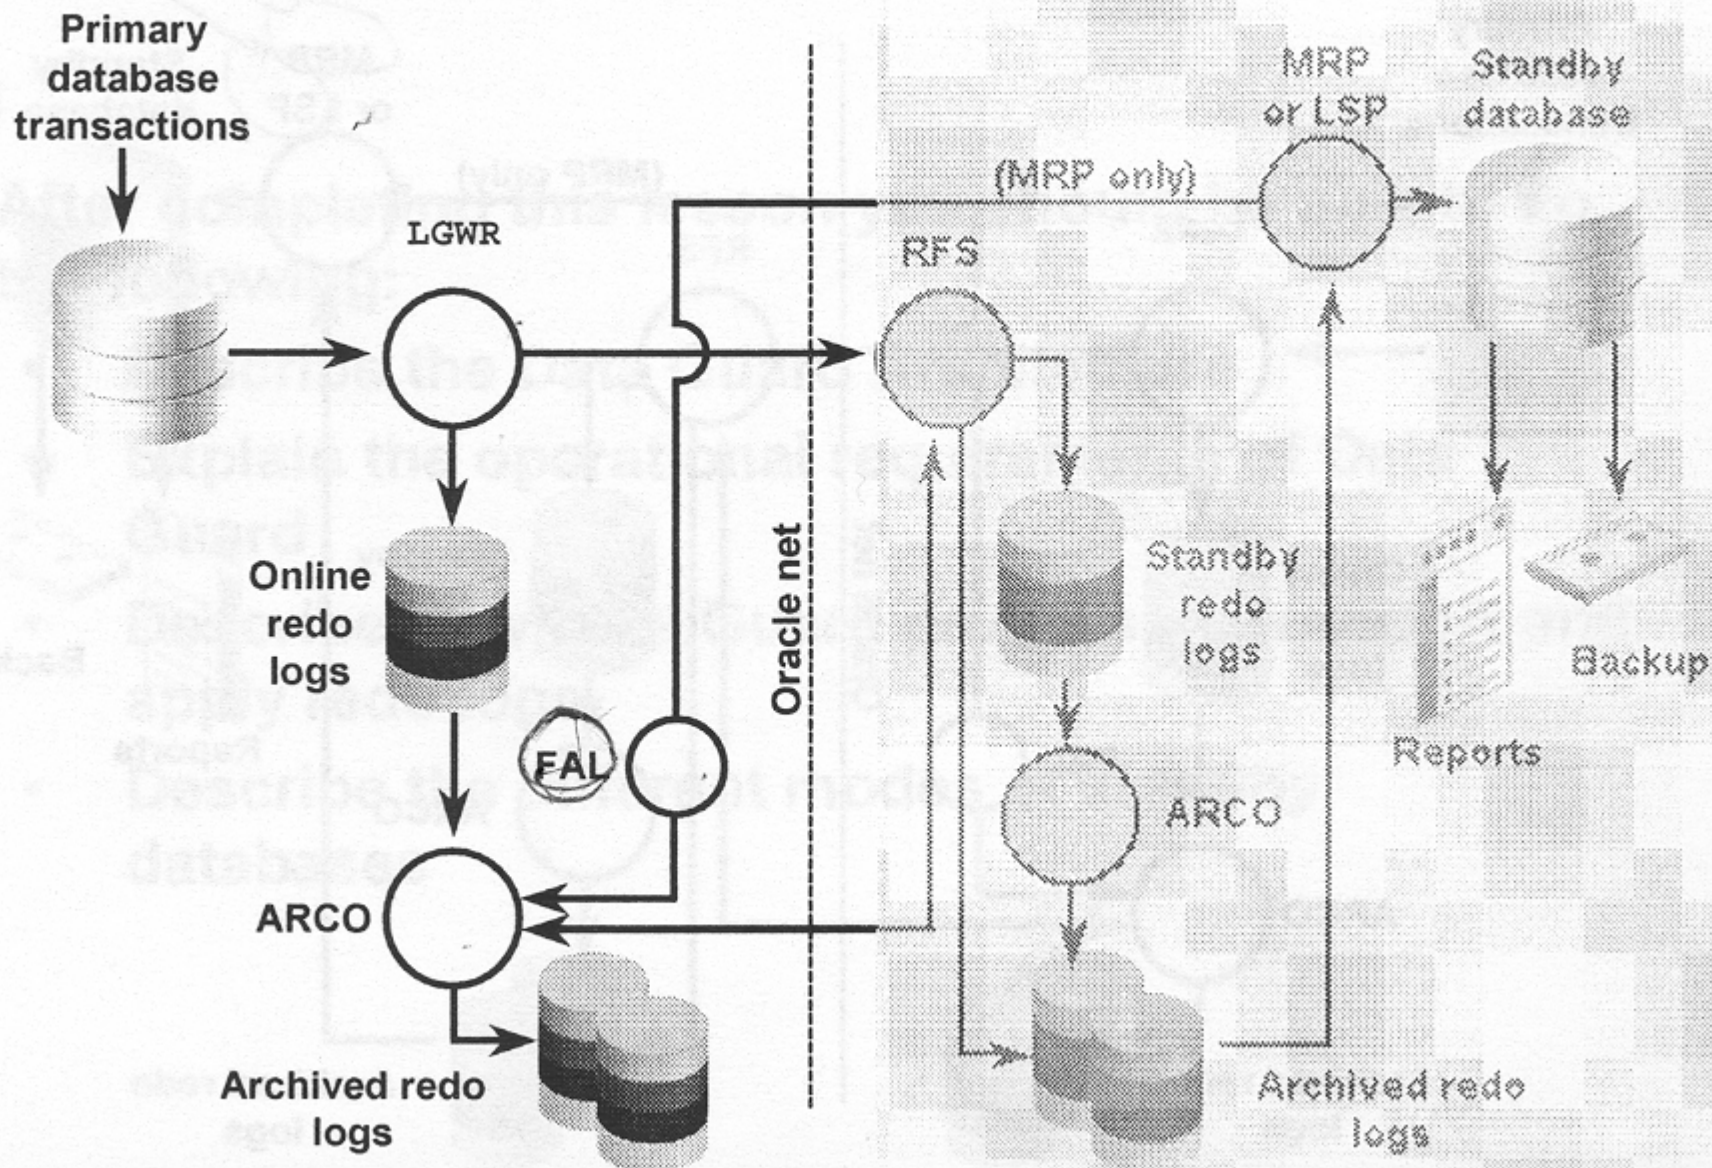
How is CDF Online protected now?

- Single Disk Failure
 - Mirroring
- Database Corruption - human error
 - Rman (potential data loss)
- Database Loss – hardware failure, major power outage(ups failure)
 - Rman
- Secondary media for point-in-time –
 - exports on demand, not regular due to lack of disk availability

What Is Oracle Data Guard?



Primary Database Flow



Where does Data Guard fit ?

- Unplanned Downtime
 - System Failure
 - Real Application Clusters
 - Data Failure and Physical Disaster
 - Data Guard
 - Human Error
 - Flashback query, rman, export/import

Types of Standby Databases

- There are two
 - Physical
 - Identical to the primary database block-for-block
 - Synchronized with the primary db thru redo apply
 - Managed Recovery State
 - Logical - not recommended with replication unsupported ddl
 - Shares the same schema definition
 - Synchronized with the primary db thru sql apply
 - Open Read-Write Mode (allows read to schemas being maintained by the redo apply and access to other db objects that are not maintained)
 - Unsupported Data Types Logical Database
 - Nclob, long, long raw, bfile, rowid and urowid, materialized view, and functional indexes

Data Protection Modes

- Maximum Protection
 - 0 data loss
 - Halts production if it can't write to standby
 - recommended at least 2 machines 2 instances standbys
- Maximum Availability
 - 0 data loss
 - Does not halt production degrades to Max Performance to catch then back to Max Availability (Data Guard Broker)
 - recommended at least 2 machines 2 instances tandbys
- Maximum Performance
 - potential data loss
 - Best performance on production
 - Only 1 standby required

Data Guard Switchover & Failover

- Neither automatically invoked, both require databases to bounce
- Switchover
 - Planned role reversal for testing
 - Read data to see if in case of disaster works
 - Used for o/s or hardware maintenance
- Failover
 - Unplanned role reversal (disaster)

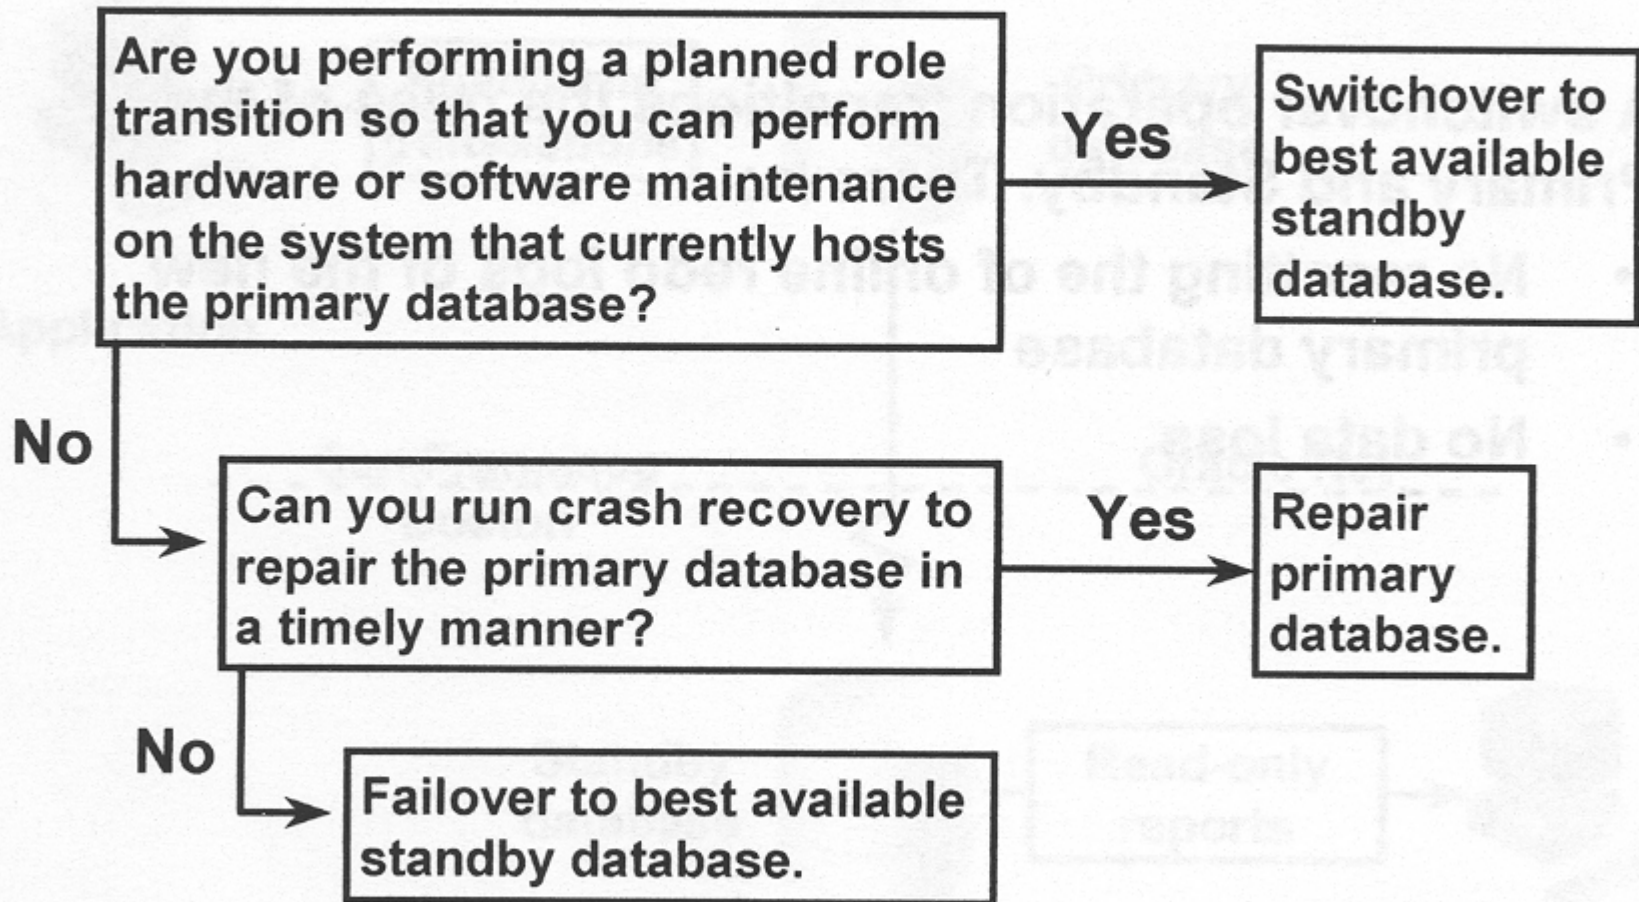
Replication in the mix

- “Replication serves two different business purposes and attempts to combine these features by using common system resources, may lead you/ the customer to uncharted/untested waters.” per Oracle

Streams in the mix

- “Generally speaking, we support a Data Guard configuration involving a primary database (A) and a physical standby database (B), along with a Streams configuration involving the same primary database (A) but a different replicated database (C). The usual Data Guard switchovers/failovers are expected to work, altho‘ some manual intervention may be required to re-establish the Streams configuration after a switchover/failover.”

Role Transition Decision Tree



Costs

- Hardware
 - Retrofit old offline machine investigated
 - Not feasible, less expensive to purchase E4500 vs retrofit
 - Likely to get a refurbished E4500 and disks
 - http://www-css.fnal.gov/dsg/internal/briefings_and_projects/planning/Data_Guard/
- High implementation, testing and maintenance costs(human) including interaction with replicated boxes

Why Data Guard?

- Continue service through a disaster with manual intervention
- Data Protection against corruption and data loss dependent on Protection Mode
- Off-load rman to your Standby
- Get more work on production system
 - Non-interruptive work is limited to additional cpu and adding disks
 - Does not include patches on oracle or o/s

Why Not Data Guard?

- Manual Intervention to failover in case of disaster
- Manual Intervention to reestablish replication
- Mix of Data Guard and Basic Replication not recommended
- Mix of Data Guard & Streams is uncharted water
- Regular testing of switchover to test standby
- Client and Application Interruption
- Additional Hardware – for 0 data loss +2
- New untested recovery scenario
 - People using backtrack product from BMC
 - Experience with RMAN not consistent

Does this meet our requirements?

- Mixed Basic Replication and Data Guard environment
- Mixed Streams and Data Guard environment
- No data loss during long downtime – required
- All aspects supported by Oracle
- Do not increase sysadmin or DBA load unreasonably
- No more than 1 run lost (4 hours data=downtime) at hard failure

Does this meet our requirements?

- Uses linux platforms that will not need replacing before the end of Run II.
- In place in less than 6 months
- < \$70K hardware
- Zero data loss at hard failure
- Transparent (no intervention) failover

Other options

- Do Nothing
 - take the downtime and recover database
- Another E4500 machine – \$10K (similar to cdfoffline)
- Pre buffer data
 - Define disk storage and load data after database recovered
- Use some other instance
 - Temporary solution to collect data then transfer over
- Investigate improving our recovery strategies
- Standby Database
 - Oracle 8 feature vs 9 simpler implementation